

Approximate Equilibrium and Incentivizing Social Coordination

Elliot Anshelevich Shreyas Sekar

Rensselaer Polytechnic Institute, Troy, NY 12180
eanshel@cs.rpi.edu, sekars@rpi.edu.

April 21, 2014

Abstract

We study techniques to incentivize self-interested agents to form socially desirable solutions in scenarios where they benefit from mutual coordination. Towards this end, we consider coordination games where agents have different intrinsic preferences but they stand to gain if others choose the same strategy as them. For non-trivial versions of our game, stable solutions like Nash Equilibrium may not exist, or may be socially inefficient even when they do exist. This motivates us to focus on designing efficient algorithms to compute (almost) stable solutions like Approximate Equilibrium that can be realized if agents are provided some additional incentives. Our results apply in many settings like adoption of new products, project selection, and group formation, where a central authority can direct agents towards a strategy but agents may defect if they have better alternatives. We show that for any given instance, we can either compute a high quality approximate equilibrium or a near-optimal solution that can be stabilized by providing small payments to some players. We then generalize our model to encompass situations where player relationships may exhibit complementarities and present an algorithm to compute an Approximate Equilibrium whose stability factor is linear in the degree of complementarity. Our results imply that a little influence is necessary in order to ensure that selfish players coordinate and form socially efficient solutions.

1 Introduction

Historically, the term coordination game has been applied to social interactions with positive network externalities. Typically, they are used to represent scenarios like the *Battle of the Sexes* wherein self-interested agents benefit if and only if they choose the same strategy. Such a model, however, does not fully capture real-life situations like the adoption of technologies or opinions and the

selection of activities where agents may eschew coordination if their personal preference for an alternative is very strong. For instance, a company may not adhere to common standards if the benefit from using their own proprietary technology far outweighs the gains from coordinating. Bearing this in mind, we consider the following broader interpretation of coordination games as ‘a class of games where agents’ utilities increase when more people choose the same strategy as them’. Notice that this does not preclude agents from having intrinsic preferences for strategies.

Many social and economic interactions fall within our framework (see [21,25] for specific applications of coordination games and different types of interactions) and it is not surprising that the kind of games we are interested in have appeared in various guises throughout literature. Researchers have studied similar kinds of games in several settings including opinion formation [12], information sharing [26], coalition formation [19] and party affiliation [7,24], largely focusing on the existence and quality of stable solutions.

Given the significance of social coordination games, a natural question that arises is: Do instances of such games result in stable outcomes that are comparable to the social optimum, the solution maximizing social welfare. The *somewhat* negative answer to this question that we provide serves as the starting point for our work as it highlights the need for incentivizing agents to form solutions that are beneficial for society. We attempt to answer the above question by articulating two fundamental drawbacks of coordination games, which naturally lead to the issue of influencing players to form “good” solutions.

1. **Coordination Failures.** Coordination games suffer from Coordination Failures [15] that result in agents becoming trapped in inefficient equilibria despite the existence of high-welfare equilibria. These situations arise when agents settle for less risky alternatives if they anticipate that other agents may not coordinate with them on what are potentially “high-risk, high-reward” solutions. As an example, consider N independent but complementary firms, each with a distinct preferred location, deciding on where to locate. Suppose that each company receives unit utility for choosing their favorite location and one more unit for every additional firm that choose to locate in the same area. Clearly an optimum and stable solution is one where all companies choose the same location. However, the outcome where each chooses their preferred location is also stable as no company could unilaterally deviate and profit. There is a large body of theoretical and experimental evidence (see [27] for a survey) that supports the hypothesis that agents may coordinate on inefficient outcomes even when better alternatives exist.
2. **Non-existence of Equilibrium.** In instances where player relationships are asymmetric (they receive different gains from coordinating), a Pure Nash Equilibrium¹ may not even exist. The following example illustrates one such simple instance with three players and three strategies.

¹We shall henceforth refer to Pure Nash Equilibrium as just Nash Equilibrium

Example There are 3 players i_1, i_2, i_3 . Player i_j receives a utility of $\sqrt{2}$ if she chooses strategy j and unit utility from strategy $j + 1$ (addition here is modulo 3). Also, player i_j receives coordination gains of 1 when player i_{j+1} chooses the same strategy as her. Note that relationships are asymmetric so player i_{j+1} receives no benefit for choosing the same strategy as player i_j .

It is not hard to verify that no Nash equilibrium exists for this instance. The reader is asked to refer to the proof of Proposition 3.1 for details.

It is evident that even in fairly simple coordination games, it may be necessary to guide agents to form desirable solutions. From a central point of view, a high social welfare is the most important requirement, but at the same time it is necessary that selfish agents do not deviate from these centrally promoted solutions. A key algorithmic challenge is therefore, computing stable outcomes with good social welfare that can be formed by providing each agent a small incentive. It is towards this end that we identify approximate equilibria as our primary solution concept.

Approximate Equilibrium and Forming Stable Solutions. An α -Approximate Equilibrium is an outcome in which no player can improve their utility by a factor more than α by unilaterally deviating. Observe from the definition that if each player is provided additional benefits equaling a fraction $(\alpha - 1)$ times their original utility, then no player would wish to deviate and the Approximate Equilibrium becomes a Nash Equilibrium. Alternatively, approximate stability corresponds to the addition of a switching cost that captures the inertia players may have in changing strategies unless the added benefit is large enough. In addition to being a simple generalization of Nash Equilibrium, approximate equilibria are also easily implementable or enforceable in natural settings as opposed to non-deterministic generalizations like Mixed Nash Equilibria.

We focus on computing approximate equilibria with high social welfare establishing that although coordination games may not admit Nash Equilibrium, the addition of a relatively small amount of inertia to the game causes stable solutions with high social welfare to exist. We also consider group deviations via approximate strong equilibrium [18], computing solutions which are resilient to deviations by sets of players. A solution \mathbf{s} is an α -Approximate Strong Equilibrium if in any deviation by a set of players from \mathbf{s} , at least one of the players' utilities cannot improve by a factor more than α as compared to her original utility. We may also refer to α as the stability factor for all approximately stable solutions.

Formalizing our Model of Coordination Games

We begin by considering a non-transferable utility game with N players and m distinct strategies. We assume that players have access to all strategies, therefore in any outcome of the game, a player's strategy $s_i \in \{1, \dots, m\}$. Generalizing our examples from the previous section, we not only permit pre-

ferred strategies, but allow each player to have asymmetric preferences over the strategy set. Formally, player i derives a utility of w_i^k if she chooses strategy $1 \leq k \leq m$.

With regards to the coordination aspect, we now propose a framework where the benefits of coordination between two players do not depend on externalities. Specifically, suppose that $w(i, j)^2$ is the total coordination benefit when players i and j choose the same strategy and that this benefit is divided among the two players. Formally, player i derives a utility of $\gamma_{ij}w(i, j)$ for coordinating with j and player j receives $\gamma_{ji}w(i, j)$. For a given instance of our game, the values $(\gamma_{ij}, \gamma_{ji}) \forall (i, j)$ are fixed and add up to one. So given a strategy vector $\mathbf{s} = (s_1, \dots, s_n)$, the utility of a player i has two components,

$$u_i(\mathbf{s}) = w_i^{s_i} + \sum_{s_j=s_i} \gamma_{ij}w(i, j).$$

We parameterize any instance on a factor γ that captures the maximum asymmetry that exists in relationships. Formally $\gamma = \max \frac{\gamma_{ij}}{\gamma_{ji}}$ over all pairs (i, j) . We term this parameter, the *Maximum Relationship Imbalance* (MRI). Since we are concerned with the quality of stable solutions, we define a social welfare function $u(\mathbf{s})$ to be the sum of player utilities. Mathematically, $u(\mathbf{s}) = \sum_i u_i(\mathbf{s}) = \sum_i w_i^{s_i} + \sum_{s_i=s_j} \gamma_{ij}w(i, j)$. We define the optimum to be the solution maximizing social welfare and compare the quality of our solutions to the optimum welfare OPT .

1.1 Our Contributions.

In this work, we consider the following well-motivated question: can we implement solutions with high social welfare by providing each player some incentive to not deviate? Our main results answer this question in the affirmative, and more importantly we show that this is possible for every instance using one of our two incentivizing schemes. First, we present an algorithm based on greedy dynamics to compute a good quality, almost-stable solution.

- (Theorem 4.5) There is a polynomial-time algorithm to compute an α -Approximate Equilibrium ($\alpha \in [1.618, 2]$) with a social welfare that is comparable to the optimum.

An approximate equilibrium corresponds to an easily realizable solution in the presence of either incentives, switching costs or players with inertia. Our second main result considers a complementary notion of stability: the minimum total payment to be provided to players so that they do not deviate from a desired high quality solution. For any given instance, if the algorithm of Theorem 4.5 returns an α -Approximate Equilibrium with social welfare $\rho_\alpha \cdot OPT$, then we show

- (Theorem 5.1) The optimum solution can be stabilized with a total payment of $\frac{\rho_\alpha}{\alpha-1}OPT$.

²We take $w(i, j)$ and $w(j, i)$ to refer to the same entity

Informally, this result tells us that if we can provide certain players supplementary utility, then with a finite budget we can stabilize the optimum solution. Given any instance, we first run the algorithm of Theorem 4.5 to compute an α -Approximate Equilibrium with social welfare $\rho_\alpha \cdot \text{OPT}$. If ρ_α is large, then we have an almost-stable solution with high social welfare; else, if ρ_α is small, then with a budget of $\approx \rho_\alpha \cdot \text{OPT}$, we can force even the optimum solution to be stable. Together our two theorems imply something much stronger: we can always *either* compute almost-stable solutions with high welfare, *or* directly stabilize a good quality solution with small payments.

In the process, we show some additional computability results for equilibria in our game.

- For $m = 2$, we can compute a Nash Equilibrium and Strong Equilibrium in polynomial time.
- For $m = 3$, we can compute a 1.414-Approximate Equilibrium and this factor is tight. i.e., for $\alpha < 1.414$, α -Approximate Equilibrium may not exist.
- For $m \geq 3$, a 1.618-Approximate Equilibrium and a 2-Approximate Strong Equilibrium always exist and we give algorithms to compute both.
- We obtain tight lower bounds for the social welfare of the solution computed by the algorithm of Theorem 4.5 in terms of γ (Maximum Relationship Imbalance). When γ is not too large, we show that this social welfare is comparable to the optimum. For instance, if relationships are not too asymmetric and $\gamma = 2$, we can compute a 2-Approximate Equilibrium whose social welfare is always at least forty percent of OPT and a 1.618-Approximate Equilibrium whose social welfare is one-third of OPT. If $\gamma < 2$, then we can do much better. Table 1 captures the social welfare of the solution returned by our algorithm for different values of m and γ .

Existence. In general, Nash Equilibrium need not exist when $m \geq 3$. However, we identify sufficient conditions that encapsulate a broad class of Social Coordination Games (even when reward sharing is asymmetric) that guarantee the existence of a Nash Equilibrium. We do this by exhibiting a novel ordinal potential function for games where the benefits of coordination among players are closely correlated across relationships. This allows the mechanism designer to identify or create settings where repeated best-response by players always converges to a stable solution. We also show that Strong Equilibrium need not exist when $m \geq 3$ even under simplifying assumptions like $\gamma_{ij} = 1 \forall (i, j)$.

1.2 Related Work.

Our model of social coordination is closely linked to two well known classes of games: non-transferable utility coalition formation and party affiliation. We begin by surveying the substantial literature in both these fields and examine the ties between our model and the games in these frameworks.

| γ | $\alpha = 2$ | | $\alpha = 1.618$ | |
|----------|--------------|------------------------|------------------|------------------------|
| | $m = 4$ | $m \rightarrow \infty$ | $m = 4$ | $m \rightarrow \infty$ |
| 1 | 0.57OPT | 0.5OPT | 0.424OPT | 0.35OPT |
| 2 | 0.47 | 0.4 | 0.37 | 0.29 |
| 10 | 0.25 | 0.15 | 0.18 | 0.12 |

Table 1: Performance of our algorithm: The social welfare of our computed solution as a fraction of the optimum welfare for different values of m (Number of strategies) and γ (Maximum Relationship Imbalance).

Hedonic games model players forming coalitions such that a player’s utility depends only on the members of her own coalition. Our model can be embedded in this setting by considering a fixed number of non-anonymous coalitions and a set of players who are anchored, i.e., constrained to join only one particular group. Much of the work in hedonic coalition formation has focused on identifying conditions for the existence of stable solutions [8, 10]. It is known, for instance, that if relationships are symmetric then the existence of Nash Equilibrium can be guaranteed by means of a potential function. Augustine et al. [3] consider a model similar to ours with the coordination benefits being submodular and characterize settings where Nash Equilibrium always exist. Although our games do not admit Nash Equilibrium, our results imply the existence of a stability concept that is slightly weaker than Nash stability for a large class of hedonic games with asymmetry.

Another line of work has focused on quantifying the inefficiency of stable solutions [11] and on the computation of stable solutions [4, 16, 20]. Although there have been a number of positive algorithmic results, the focus on approximating both stability and optimality has been limited. We also remark here that many of the defection models considered in the hedonic games literature are relevant for settings with an unbounded number of anonymous coalitions and may not be suitable for our class of games. With regards to influencing, the recent work on stabilizing desired coalitions via supplementary payments (the Cost of Stability) [5] is similar to our direct payments technique, albeit in a transferable utility setting.

Party affiliation games are a generalization of pure coordination games where players wish to coordinate with friends and *anti-coordinate* with enemies. On the other hand, in our model the friction is provided by the interplay between a player’s individual preference and coordination. Party affiliation games with only two strategies and symmetric relationships have received considerable attention as Nash Equilibrium always exists in these settings although computing it is PLS-Complete [13]. As a positive algorithmic result, Bhalgat et al. ([9]) gave a polynomial time algorithm to compute a $(3 + \epsilon)$ -Approximate Equilibrium for such games. Most similar to our work in terms of motivation is the paper by Balcan et al. ([7]) who consider a model where only a limited number of players follow the centrally advertised strategy. Even though we look at only one aspect of party affiliation, our model is quite general as we do not impose

any restriction on the number of strategies or player relationships.

Other models of coordination in strategic settings. There has been a renewed interest in studying coordination games in networks from a theoretical perspective with an emphasis on identifying the kind of equilibrium outcomes that emerge. In particular, work has focused on characterizing the effect of several parameters on the equilibrium outcome including the cost of forming links [22, 25], network structure [14], level of interaction [32] and incomplete information [21]. The literature on coordination games is too vast and the reader is asked to refer to the paper by [21] for a more detailed survey.

2 Preliminaries and Warm-up Results

In this section, we address some fundamental questions regarding stability and optimality in our Social Coordination Game (SCG) in order to gain a better understanding of our model. We then move on to the most simple version of our game where players only have two different strategies to choose from. For this special case, we present an algorithm that always returns a Nash Equilibrium whose social welfare is half of the optimum. We begin by casting our game in graph theoretic framework to compare our problem to existing optimization problems.

Social Coordination as a Network Game. We can view our model as a game played on a complete graph $G = (V, E)$ where the nodes include the players and m additional anchored nodes (which are constrained to choose only one strategy). Each directed edge (i, j) has a weight $\gamma_{ij}w(i, j)$, the utility player i derives from coordinating with j . Recall that player j derives a utility of $\gamma_{ji}w(i, j)$ from coordinating with player i and a utility of w_j^k for choosing strategy $1 \leq k \leq m$.

We now define some additional notation to understand the social welfare of solutions better. We divide the social welfare of any strategy vector \mathbf{s} into two components. $u(\mathbf{s}) = A(\mathbf{s}) + P(\mathbf{s})$, where $A(\mathbf{s})$ refers to the utility players receive due to their intrinsic preference for the chosen strategy and $P(\mathbf{s})$, the contribution of players' relationships to social welfare, both under \mathbf{s} . Mathematically $A(\mathbf{s}) = \sum_i w_i^{s_i}$ and $P(\mathbf{s}) = \sum_{s_i=s_j} w(i, j)$. We denote by $best(i)$, the maximum utility that player i derives from any one strategy and use A_T to express $\sum_i best(i)$. Similarly, we use P_T to refer to $\sum_{(i,j)} w(i, j)$, the maximum possible contribution of players' relationships to social welfare. Finally, we can upper bound OPT by $A_T + P_T$.

It is not hard to see that the problem of maximizing social welfare is equivalent to the problem of dividing the graph into m clusters to maximize the weight inside the clusters. This in turn is equivalent to the problem of minimizing the weight of the edges going across different clusters, which is the popular *Multiway Cut* problem. We infer therefore, that the problem of maximizing social welfare in any given instance of the Social Coordination Game is NP-Hard.

Proposition 2.1. *For $m > 2$, computing the optimum solution for an instance of the Social Coordination Game is NP-Hard even for a fixed value of γ .*

The optimization version of our problem was considered for undirected graphs in [28] as the *Multiway Uncut problem*. They exhibited a 0.8535 approximation algorithm for the same using an LP-rounding based approach. We extend their results to the directed version via a simple reduction.

Proposition 2.2. *There exists a polynomial-time algorithm to compute a solution of the SCG such that its social welfare is at least a fraction 0.8535 of OPT .*

Proof. Given an instance of the SCG, we create an undirected graph with weight $w(i, j)$ on edge (i, j) and run the algorithm of Langberg et al. \square

Social Coordination Games with Two strategies

The simplest form of Social Coordination is one where players get to choose between two distinct strategies, which has received considerable attention in literature. For this special case, we are able to guarantee the existence of a Nash Equilibrium by constructing one for any given instance. Our actual result is much stronger, we give a sequence of best-response moves from any given starting state that always results in a Nash Equilibrium. By Proposition 3.3 (shown in the next section), we are guaranteed that this solution is at least half as good as OPT in terms of social welfare.

Algorithm 1 Computing Nash Equilibrium when $m = 2$

- 1: Allow players to deviate from strategy 1 to 2 if it is their best response until no player wants to deviate from strategy 1.
 - 2: Repeat this for the other strategy until no player can improve her utility by deviating from strategy 2.
-

Proposition 2.3. *The above algorithm returns a Nash Equilibrium from any given starting state \mathbf{s}_0 such that its social welfare is at least half of OPT .*

Proof. At the end of Step 1, no player would want to deviate from strategy 1 to 2 by definition. These players would not want to deviate after step 2 either as their utility only increases. By definition, players cannot improve their utility by deviating from strategy 2 and thereby we have a Nash Equilibrium. \square

A similar algorithm was used in [12] to compute Nash equilibria in a cost minimization game with two preferences.

3 Existence and Quality of Stable Solutions

In the previous section, we showed that if every player has exactly two strategies to choose from, we can always compute an equilibrium in polynomial time. Unfortunately, we are not so lucky even if players have one additional strategy to choose from as Nash Equilibrium may not exist. However, in this section,

we show sufficient conditions to guarantee the existence of equilibrium when $m \geq 3$. In the context of our central theme of incentivizing coordination, this is useful in settings where the central designer may not be able to provide any direct incentives but can exert some control over the parameters of the game so that natural game play (best-response dynamics) results in stable solutions.

Although similar games like Party affiliation and Discrete opinion formation also admit potential functions [9, 12], our result does not follow from theirs as our potential function is not an exact potential function. Further, our conditions also capture scenarios where the benefits of coordination are not equally shared. First, we formalize the example given in the introduction regarding the non-existence of equilibrium.

Proposition 3.1. *There exist instances of the Social Coordination Game with three strategies where Nash Equilibria do not exist.*




Figure 1: Instance of the SCG with 3 strategies (represented by shaded anchor nodes) and three players (i_1, i_2, i_3). Player i_j receives a utility of $\sqrt{2}r$ if she chooses strategy x_j and r utility from strategy x_{j+1} . Also, player i_j receives coordination gains of r when player i_{j+1} chooses the same strategy as her.

Proof. In any equilibrium, each player receives a utility of at least $\sqrt{2}r$ or else they could deviate to their favorite strategy. It is not hard to see that in any strategy other than the one where all players are separated, at least one player's utility is strictly less than $\sqrt{2}r$ and hence she would deviate. Therefore, in every strategy at least one player can deviate and improve her utility by a factor of $\sqrt{2}$. \square

Although existence cannot be guaranteed for the general case, the following natural assumption captures a broad class of SCGs for which Nash Equilibrium always exists.

(Correlated Coordination Condition) A given instance of the SCG is said to satisfy this condition if \exists a vector $\gamma = (\gamma_1, \dots, \gamma_N)$ such that $\forall (i, j)$ with $w(i, j) > 0$, we have $\gamma_{ij} = \frac{\gamma_i}{\gamma_i + \gamma_j}$ and $\gamma_{ji} = \frac{\gamma_j}{\gamma_i + \gamma_j}$.

Informally, each player has an intrinsic weight or an influence factor γ_i . When two players (say i and j) coordinate, the benefit that any one player

derives depends on her relative weight with respect to the other player. For any two players i and j , if $\gamma_i \geq \gamma_j$, then player i receives a greater benefit due to coordinating with player j . Notice that this definition does not impose any restriction on how asymmetric relationships can be. We now claim that games which obey this condition not only have Nash Equilibrium but in fact admit an ordinal (inexact) potential function [31] that ensures that best-response always converges to an equilibrium.

Theorem 3.2. *Social Coordination Games which obey the Correlated Coordination (CC) Condition admit a potential function. Therefore, best-response dynamics always converge to a Nash Equilibrium in such games.*

Proof. Given a social coordination game with weights $w(i, j), \forall (i, j)$ and $\gamma = (\gamma_1, \dots, \gamma_N)$, that obeys the CC condition, we claim that the following is an inexact potential function for this game.

$$\Phi(\mathbf{s}) = \sum_i \frac{w_i^{s_i}}{\gamma_i} + \sum_{s_i=s_j} \frac{w(i, j)}{\gamma_i + \gamma_j}.$$

To prove this, we need to show that whenever a player makes an improving move, the potential function increases and vice-versa. Consider a strategy vector \mathbf{s} and a player i whose utility increases, when she deviates from s_i to s'_i . The change in her utility is $u_i(s'_i, s_{-i}) - u_i(s_i, s_{-i}) = (w_i^{s'_i} + \sum_{s_j=s'_i} \gamma_{ij} w(i, j)) - (w_i^{s_i} + \sum_{s_j=s_i} \gamma_{ij} w(i, j))$, which is positive. As our games obey the CC condition, we can substitute γ_{ij} with $\frac{\gamma_i}{\gamma_i + \gamma_j}$ and obtain,

$$w_i^{s'_i} + \sum_{s_j=s'_i} \gamma_i \frac{w(i, j)}{\gamma_i + \gamma_j} - w_i^{s_i} - \sum_{s_j=s_i} \gamma_i \frac{w(i, j)}{\gamma_i + \gamma_j} > 0.$$

This implies that,

$$\gamma_i \left(\frac{w_i^{s'_i}}{\gamma_i} + \sum_{s_j=s'_i} \frac{w(i, j)}{\gamma_i + \gamma_j} - \frac{w_i^{s_i}}{\gamma_i} - \sum_{s_j=s_i} \frac{w(i, j)}{\gamma_i + \gamma_j} \right) > 0.$$

Since γ_i is always positive, we conclude that the term inside the parenthesis must be positive as well. Therefore, the potential function strictly increases when players make a better-response move. The other direction is similar, if $(\frac{w_i^{s'_i}}{\gamma_i} + \sum_{s_j=s'_i} \frac{w(i, j)}{\gamma_i + \gamma_j} - \frac{w_i^{s_i}}{\gamma_i} - \sum_{s_j=s_i} \frac{w(i, j)}{\gamma_i + \gamma_j}) > 0$, then we multiply both sides by γ_i yielding the change in player i 's utility on the left hand side. This indicates that any deviation by a player which increases the potential function must be a strictly improving move. We conclude that these games admit a potential function. Therefore, best-response dynamics must always converge to a Nash Equilibrium. \square

Among other scenarios of interest, the CC conditions also captures games where player relationships are symmetric but players may derive any arbitrary utility from choosing a strategy ($\gamma_i = \gamma_j \forall (i, j)$). Therefore, if the benefits of a relationship are split equally among the players, stable solutions always exist.

Quality of Stable Solutions

Although Nash Equilibrium exists for special cases, it is not clear whether these stable solutions would be centrally desirable in terms of social welfare. We now formalize the notion of “Coordination Failures” that we mentioned earlier and show that in general Coordination Games, both the best and worst Nash Equilibrium can be a factor m away from the social optimum. In other words, even when Nash Equilibrium exists, all stable solutions may have a social welfare that is only a small fraction of OPT . This reinforces the need for some external influence in order to ensure solutions that are closer to OPT in terms of social welfare.

Proposition 3.3. *The Price of Anarchy (PoA) for Social Coordination Games over all instances is at most m .*

The Price of Anarchy is the ratio of the social welfare of the optimum to that of the worst Nash Equilibrium (equilibrium with the lowest social welfare for a given instance) over all instances.

(*Proof Sketch*) We use a refinement of the popular smoothness technique [33] known as semi-smoothness (see [2, 30] for more details) for proving the above bound. The full proof is present in the appendix. We show that our Social Coordination Game is $(\frac{1}{m}, 0)$ -semi-smooth for all instances. This immediately gives us an upper bound of m for the Price of Anarchy. ■

Note that by the generality of the smoothness technique, we are guaranteed that the Price of Anarchy is m for solution concepts that are much more general than Nash Equilibrium including Mixed Nash Equilibrium and Coarse Correlated Equilibrium. We now show that the above factor is tight by giving a matching lower bound for the best Nash Equilibrium.

Proposition 3.4. *The Price of Stability (PoS) for the SCG is at least m and this is tight.*

The Price of Stability for a given game is the maximum ratio of the social welfare of the optimum and the best Nash Equilibrium over all instances.

Proof. (Example) Consider an instance with m strategies and m players, $i_1 \cdots i_m$. Player i_1 derives a utility r from strategy i and an additional utility of r for every other player who chooses the same strategy as her. Players i_j ($j > 1$) receive a utility of 2ϵ from choosing strategy j and an additional utility ϵ if player i_1 chooses the same strategy. In the social optimum, all players choose strategy 1, and we have $u(OPT) = mr + (m - 1)\epsilon$. This instance has m different Nash Equilibria based on which strategy player i_1 chooses and every Nash Equilibrium has a social welfare of $r + 2(m - 1)\epsilon$. In the limiting case as $\epsilon \rightarrow 0$, the ratio between the social welfares approaches m . □

This implies that in instances where Nash Equilibrium exist, all stable solutions can have a social welfare that is much worse than the optimum (if m is large). Propositions 3.3, 3.4 also tell us that if we are able to limit the number of

choices available to self-interested agents, then all stable solutions exhibit high social welfare.

Quality of Stable Solutions for the Symmetric Coordination Game

We return to the special case when player relationships are symmetric (rewards are split-up equally), which is of considerable interest as it occurs in several scenarios. Notice that this case can be expressed as a SCG with $\gamma_{ij} = 1 \forall (i, j)$ and therefore, every player receives a utility of $\frac{w(i,j)}{2}$ from each relationship. We show that symmetry improves the situation by giving a tight upper bound of $(2 - \frac{1}{m})$ on the Price of Stability. Unfortunately, this best equilibrium may not be computable. In the next section, we present an algorithm to compute an approximate equilibrium that is almost as good as the best Nash Equilibrium. First observe that the following is an exact potential function for the SCG with symmetric relationships.

$$\Phi(s) = \sum_i w_i^{s_i} + \frac{1}{2} \sum_{s_i=s_j} w(i, j).$$

Proposition 3.5. *The price of stability is at most $(2 - \frac{1}{m})$ in social coordination games with symmetric relationships, i.e. $\gamma = 1$. Further, this bound is tight.*

(*Proof Sketch*) The complete proof can be found in the Appendix. We prove the first part of the proposition by proving a much stronger statement on best-response dynamics: that the loss in social welfare from certain ‘good starting states’ after any best-response sequence is upper bounded by the factor $2 - \frac{1}{m}$. Then all we need to show is that OPT belongs to these good states, by proving that $A(OPT) \geq \frac{Ax}{m}$. We also show a tight example where the PoS is exactly $(2 - \frac{1}{m})$, which completes the proof. ■

The proposition also implies that in the special case with two strategies, the polynomial-time best-response sequence described previously with OPT as a starting state (can be computed using any algorithm for Min-Cut) leads to a Nash equilibrium whose social welfare is at least two-thirds of OPT .

Corollary 3.6. *Algorithm 1 for two strategy SCGs with symmetric relationships returns a Nash Equilibrium that is at least two-thirds of the optimum when OPT is used as the starting state.*

4 Approximate Equilibrium

We now come to our main results. In the previous sections, we showed that Nash Equilibrium may not exist for Social Coordination Games and even when it does, its quality may not be comparable to the optimum. This motivates us to relax our notion of stability and consider approximately stable solutions in order to obtain better guarantees on the social welfare. Our first main result in this section is a polynomial-time algorithm to compute a 1.414-Approximate

Equilibrium for SCGs with three strategies. When $m > 3$, we give an easily implementable algorithm to compute an α -Approximate Equilibrium for $\alpha \in [\phi, 2]$ ($\phi \approx 1.618$). As we showed that α -Approximate Equilibria for $\alpha < 1.414$ do not exist for general SCGs, our stability results are nearly tight (up to a factor of ≈ 1.14). Recall that in an α -Approximate Equilibrium, no player can deviate from one strategy to another and increase her utility by more than α times her previous utility.

Three Strategies: For $m = 3$, recall that in the example of Figure 3, in every state there is at least one node that can deviate and improve its utility by a factor of $\sqrt{2}$. We show now that this bound of $\sqrt{2}$ is tight by exhibiting an algorithm for three groups that always returns a $\sqrt{2}$ -Approximate Nash equilibrium.

Proposition 4.1. *The following algorithm returns a $\sqrt{2}$ -Approximate Equilibrium for all instances with three strategies and no algorithm can give a better approximation factor over all instances as this ratio of 1.414 is tight.*

1. Run *Algorithm 1* for $m = 2$ (for say strategies 1 and 2) and all players so that no player wants to deviate from one of those two strategies to the other.
2. While there exists a player, who can deviate from either strategy 1 or strategy 2 and improve her utility by a factor $\sqrt{2}$ or more, allow her to deviate to strategy 3. Once that player deviates, run the algorithm for $m = 2$ (for strategies 1 and 2) so that players choosing these two strategies are stable with respect to each other.
3. Repeat step 2 until no player can to deviate to 3 and increase her utility by a factor $\sqrt{2}$.

In other words, at every stage of the algorithm, players in strategies 1 and 2 are stable with respect to each other. Players deviate to 3 only when it is their best-response and it improves their utility by the given factor. The algorithm terminates in polynomial-time as in each round the size of group 3 increases by one. We prove that its stability factor is $\sqrt{2}$ in the appendix.

Moving on to the general case, we give an algorithm to compute a α -Approximate Equilibrium and characterize the trade-off between stability and optimality by showing how the social welfare increases when we decrease the stability factor from $\alpha = 2$ to $\alpha = \phi$. Figure 2(a) illustrates this near-linear trade-off between social welfare and α . As mentioned, our bounds for social welfare depend on γ , the Maximum Relationship Imbalance (MRI), that measures the maximum asymmetry over all relationships. We explicitly use this factor γ in our welfare as often real-life relationships are bounded in their asymmetry; if one player receives large gains from coordinating with another, then the second player derives at least some fraction of that benefit. Recall that $\gamma = 1$ denotes the case with symmetric relationships, as γ increases, reward sharing becomes more and more asymmetric.

We begin by presenting a simple algorithm called *one-shot α -BR*, which allows each player to play her best-response at most once as long as the improvement in utility is by a multiplicative factor $\alpha \geq 1$. This natural algorithm has desirable properties in terms of both stability and welfare and acts as a subroutine for many of our results in this section and the next. First, we show the following lemma.

Algorithm 2 One-shot α -BR

Require: A starting strategy k_0 .

- 1: Set all players' initial strategy to be k_0 .
 - 2: While there exists a player whose current strategy is k_0 , who can improve her utility by at least a factor α by deviating to another strategy, allow her to perform best-response.
 - 3: Return the final solution \mathbf{s} .
-

Lemma 4.2. *The One-shot α -BR algorithm returns a $(\max(\alpha, \frac{1}{\alpha} + 1))$ -approximate equilibrium, for any starting strategy k_0 .*

That is, for $\alpha \in [\phi, 2]$, the *One-Shot α -BR* algorithm returns an α -Approximate Equilibrium.

Proof. We introduce some notation here in order to analyze our *one-shot α -BR algorithm*. First, assume that the algorithm proceeds in steps, such that at each time step only one player deviates from k_0 . Clearly, the algorithm terminates after at most N time steps. Let X_k^t denote the set of players whose strategy is k at time step t and s_{-i}^t , the strategy vectors of all players other than i at that time. Let $X_k(\mathbf{s})$ refer to the players choosing strategy k in the final solution. Our first observation is that the number of players choosing a particular strategy other than k_0 cannot decrease with time.

Now, consider any player i such that i deviated to some strategy $k = s_i$ at time t . Then, we have

$$\begin{aligned} u_i(s_i, s_{-i}^t) &\geq u_i(k', s_{-i}^t) \quad \forall k' \neq s_i \\ u_i(s_i, s_{-i}^t) &\geq \alpha u_i(k_0, s_{-i}^t) \end{aligned} \tag{1}$$

Since i performs her best-response, she derives more utility from her strategy than she would by deviating to any other strategy. Each player who deviates improves her utility by a factor α by definition, which gives us the second inequality above. Now suppose that i 's best-response under \mathbf{s} is to move to some strategy k' . Observe that $X_{k'}(\mathbf{s}) \subseteq X_{k'}^t \cup X_{k_0}^t$ as the new players choosing k' after time t can only come from the k_0 . The utility that i gets from this best-response strategy k' can be upper bounded as $u_i(k', \mathbf{s}_{-i}) \leq u_i(k', s_{-i}^t) + u_i(k_0, s_{-i}^t) \leq u_i(s_i, s_{-i}^t) + \frac{u_i(s_i, s_{-i}^t)}{\alpha} \leq (1 + \frac{1}{\alpha})u_i(\mathbf{s})$. The last two inequalities come from invariant 1 and the fact that $X_k^t \subseteq X_k(\mathbf{s})$ respectively, giving us the stability factor for players who deviated from k_0 .

By definition, the algorithm terminates when no player whose strategy is k_0 can deviate and improve her utility by a factor α . Therefore, for any player $i \in X_{k_0}(\mathbf{s})$ and other strategy $k' \neq k_0$, we have $\alpha u_i(\mathbf{s}) \geq u_i(k', s_{-i})$. This means that any arbitrary player in N can perform her best-response and improve her utility by a factor no greater than $\max(\alpha, \frac{1}{\alpha} + 1)$, which completes the proof. \square

The following lemmas give us an idea about the social welfare of the solution returned by the *One-Shot α -BR* algorithm. In particular, we see that players get at least as much utility as they would by choosing their favorite strategy (in the absence of coordination). Further, the final utility is quite close to the utility of the starting state where all players choose the same strategy k_0 . Recall that A_T is the total utility received by players from their preferred strategy, and $A(\mathbf{s})$ and $P(\mathbf{s})$ are the utilities due to intrinsic preference and coordination in any given strategy \mathbf{s} .

Lemma 4.3. *The 1-shot α -BR algorithm, beginning with any starting strategy k_0 , returns a solution whose social welfare is at least $\frac{A_T}{\alpha}$*

Lemma 4.4. *The one-shot α -BR algorithm with \mathbf{s}_0 as the starting state (where all players choose a strategy k_0), results in a social welfare of at least $A(\mathbf{s}_0) + \frac{\alpha}{\gamma+1}P(\mathbf{s}_0)$, where γ is the Maximum Relationship imbalance.*

(*Proof Sketch*) The full proofs can be found in the appendix. Notice that the algorithm allows each player to deviate exactly once, so each player gets as much as she would (discounted by α) by choosing her preferred strategy. Summing up over all players gives us Lemma 4.3. For the second lemma, observe that if a player i deviates, and the utility of other players drops by some amount y , then player i 's original utility was at least $\frac{1}{\gamma}y$. Since she performs a best response move, her new utility is at least the same amount. Therefore, for every $(\gamma + 1)$ units of social welfare lost due to coordination, at least one unit of social welfare is gained, giving us the desired bound. \blacksquare

We are now in a position to show our main result. Since the above algorithm returns a solution whose social welfare is close to that of the starting state, it seems to natural to select a starting state with a high social welfare. We choose the strategy k^* that players have the maximum intrinsic preference for, i.e., $k^* = \arg \max_{1 \leq k \leq m} (\sum_i w_i^k)$.

Theorem 4.5. *The following algorithm returns an α -Approximate Equilibrium for $\alpha \in [\phi, 2]$ whose social welfare is approximately at least a fraction $\approx \max(\frac{\alpha}{\gamma+3}, \frac{1}{m})$ of the optimum social welfare.*

Algorithm: “For a given α , run *One-shot α -BR* and *One-shot $\frac{1}{\alpha-1}$ -BR* with k^* as the starting strategy. Let the returned solutions be \mathbf{s}_1 and \mathbf{s}_2 . Return the solution among these two with greater social welfare.”

Discussion. The exact social welfare is: $\frac{\alpha - 1}{1 + \frac{(\gamma+1)}{\alpha}(\alpha - (1 + \frac{1}{m}))}$ times OPT when $(\gamma + 1) \leq \alpha m$, and $\max(\frac{\alpha}{\gamma+1}, \frac{\alpha-1}{(m-1)+(\alpha-1)})$ times OPT otherwise. We attempt to break-down the dependence of social welfare on different parameters. First, notice from the approximate formula $\frac{\alpha}{\gamma+1}$ that the social welfare increases almost linearly with α . In other words, as is usually the case, sacrificing a little individual stability results in greater overall well being. At the same time, we see that social welfare decreases when γ becomes larger, i.e., when relationships become more asymmetric. Therefore, for a designer, there are two possible measures to ensure socially efficient outcomes: **(i)** imposing a higher switching cost on agents, **(ii)** splitting the rewards of coordination almost equally.

Our algorithm provides good guarantees when γ is not too large. For instance, when the benefits from coordination are off by at most a factor of 2 ($\gamma \leq 2$), our algorithm returns a 1.618-Approximate Equilibrium that is almost one third of OPT even if the number of available alternatives is arbitrarily large. Contrast this with the result in Proposition 3.3 that states as $m \rightarrow \infty$, the social welfare of equilibria becomes infinitely worse off than the optimum. For the same γ , if we sacrifice some stability and move to a 2-Approximate Equilibrium, we get more than forty percent of the optimum in terms of welfare. Finally, the social welfare is bounded by $\approx \frac{1}{m}$, which indicates that even when γ is large, our solutions are at least as good as that of the Nash Equilibrium (when they exist). In other words, running our algorithm is always preferable to letting agents form their own strategies. Figure 2(b) illustrates how the social welfare drops when relationships become more asymmetric.

(Proof Sketch) Both the procedures in the algorithm result in an α -Approximate Equilibrium. Applying Lemma 4.2 to the solution returned by the 1-shot $\frac{1}{\alpha-1}$ -BR procedure, we get that its stability factor is $\max(\frac{1}{\alpha-1}, \frac{1}{\frac{1}{\alpha-1}} + 1)$, that is $\max(\frac{1}{\alpha-1}, \alpha)$. \mathbf{s}_1 is, therefore, an α -approximate Nash equilibrium as $\alpha \geq \frac{1}{\alpha-1}$ for $\alpha \in [\phi, 2]$. Similarly, applying the lemma to \mathbf{s}_2 , we get that its stability factor is $\max(\alpha, \frac{1}{\alpha} + 1)$, which is also α for the given range. This establishes the stability factor.

The social welfare of the state where all players have chosen strategy k^* is at least $\frac{A_T}{m} + P_T$ as players have greater overall intrinsic preference for this strategy than any other strategy. Applying Lemma 4.4 to our algorithm's α -BR procedure, we get that the final social welfare of the solution, $u(\mathbf{s})$, is at least $\frac{A_T}{m} + \frac{\alpha}{\gamma+1}P_T$. Applying Lemma 4.3 to our One-shot $\frac{1}{\alpha-1}$ -BR algorithm, we get that our computed solution has a social welfare of at least $(\alpha - 1)A_T$. Taking the best solution among the two and finding the point where it is minimum gives us the desired bound. ■.

We return to the case where $\gamma = 1$. We previously proved that the best Nash Equilibrium has a social welfare which is a factor $2 - \frac{1}{m}$ less than OPT . Substituting, $\gamma = 1$ in the above equation, we see our 2-Approximate Equilibrium has the same social welfare as the best Nash Equilibrium.

Corollary 4.6. *For the special case when coordination benefits are split equally*

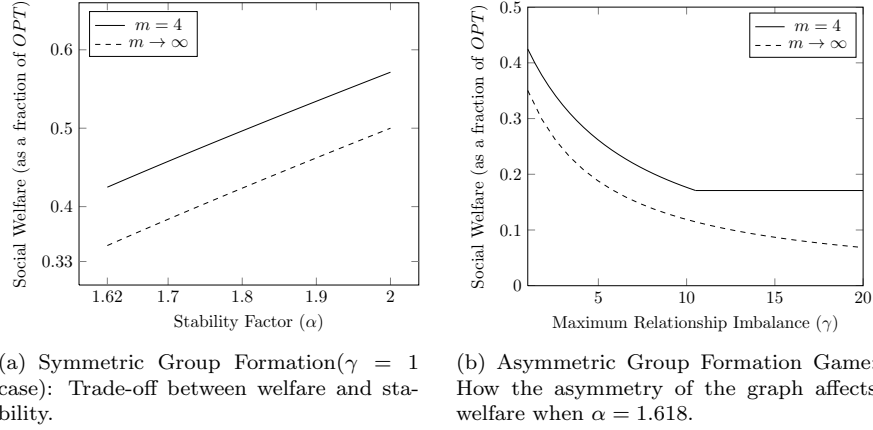


Figure 2: Social Welfare as a function of α and γ .

($\gamma = 1$), our algorithm returns a 2-Approximate Equilibrium that is more half as good as the optimum and a 1.618-Approximate Equilibrium that is at least a fraction 0.35 of OPT.

4.1 Strong Equilibrium

In any game involving coordination, it becomes important to consider Strong Nash stability as coalitions of players may be able to cooperate when switching to different strategies. Given our focus on computing good, approximately stable solutions, it becomes imperative to design solutions where for any group of players who can deviate and improve their utility, at least one player's utility increases by a factor non more than α . We term this α -approximate Strong equilibrium. In this section, we show both (non-)existence and computability results for Strong and approximately Strong equilibria.

We begin by looking at Social Coordination games with only two strategies. Our first result is that any SCG with two strategies admits a Strong Nash Equilibrium. We give a constructive proof for this by extending Algorithm 1 from Section 2 to also compute SNE.

Proposition 4.7. *Any instance with $m = 2$ strategies admits a Strong Nash equilibrium. The following algorithm computes a Strong Nash Equilibrium.*

1. Let \mathbf{s}_0 be the starting state with all players present in 1.
2. While there exist any group of players who can deviate from strategy 1 to strategy 2 such that each player's utility increases strictly from their previous utility, allow this group to deviate.

Proof. The proof is similar to that of Algorithm 1. Observe that by definition, no set of players can deviate from 1 and improve their utility. If a set of players

from 2 deviate to 1, then it is not hard to see that the player in this set who deviated at the earliest time step in the algorithm cannot be improving her utility. \square

For three or more strategies, strong equilibrium may not exist even for the simplest special case of Symmetric SCGs where $\gamma = 1$ and Nash Equilibrium is guaranteed to exist. The non-existence of strong equilibrium motivates us to look at approximate Strong equilibria and how to compute them efficiently. Our main result in this section is the efficient computation of a 2-approximate Strong equilibrium with good social welfare. First, we exhibit a non-trivial example to prove the non-existence of Strong Nash Equilibrium.

Proposition 4.8. *There are instances of the symmetric Social Coordination Game with $m = 3$ groups where Strong equilibrium does not exist. However, whenever Strong equilibrium exists, its social welfare is no less than half of OPT when player relationships are symmetric.*

Proof. Consider the instance shown in Figure 3 with 3 strategies and 5 players. The figure makes use of the graph theoretic framework defined in 2. As the instance is symmetric, we assume that utility recieved due to coordination by both the players equals the weight on the edge ($\frac{w(i,j)}{2}$). In any Nash equilibrium for this instance, two players from two different strategies can always deviate to the third strategy and improve their utilities. For instance, consider the state shown in the figure, i_2 and i_5 can deviate to strategy 2 and both increase their utilities. In any state where i_2 is not present in strategy 1, i_1 and i_3 can jointly deviate to 3. In all other states, at least one player can perform best-response by herself. We prove the second part of the claim in the appendix. \square

Proposition 4.9. *For $\alpha \geq 1$, the One-shot α -BR returns a $(1+\alpha)$ -Approximate Strong Equilibrium.*

We remark here that for $\alpha = 1$, we get a 2-Approximate Strong Equilibrium whose social welfare is at least half of that mentioned in Theorem 4.5. Therefore, our *One-shot α -BR* algorithm also offers resilience against group deviations. The full proof of the above proposition is present in the Appendix.

(Proof Sketch) Let \mathbf{s} be the solution returned. For any player i whose strategy is k_0 , for any other strategy k , we have $\alpha u_i(\mathbf{s}) \geq u_i(k, \mathbf{s}_{-i})$. Therefore, after any group deviation to strategy k (if the final strategy is \mathbf{s}'), we have $u_i(\mathbf{s}') \leq u_i(k, \mathbf{s}_{-i}) + u_i(\mathbf{s}) \leq (\alpha + 1)u_i(\mathbf{s})$, which gives us the bound. For players whose strategy is other than k_0 , their utility cannot increase a factor more than $1 + \frac{1}{\alpha} \leq (1 + \alpha)$. \blacksquare

5 Incentivizing Players

So far, we have looked at approximate equilibria for our game, which become fully stable when there are some incentives provided to players or when there is




Figure 3: Example: Symmetric Social Coordination Game with three strategies marked by ellipses (anchored players are shaded) and five players. For the purpose of clarity, edges between players are solid, whereas the edges between players and anchors are dashed. In the current state, all players within an ellipse have their strategies set to be that corresponding strategy.

limited inertia in switching. We now consider a more explicit incentivizing technique where a central authority can provide arbitrary payments to any subset of players in order to enforce a desired solution. Ideally, this desired solution is the social optimum although we later consider other high welfare solutions which are easily computable. In such situations, there are often budgets constraints which determine the total such payments that can be made and we wish to obtain bounds on the minimum budget required for any given instance. Our main result here is the complement of Theorem 4.5. Informally our result implies that when the solution returned by the Algorithm in Theorem 4.5 does not have a high social welfare, then we can ‘stabilize’ a solution that is close to OPT with a very small budget. Together, the results indicate that for any instance, we can either compute a good solution or provide small payments to enforce a solution close to OPT .

We now define our incentive scheme: we are interested in ‘stabilizing’ a given solution \mathbf{s} by providing each player i a payment of $\nu_i OPT$, where $0 \leq \nu_i \leq 1$, such that her total utility is now $u_i(\mathbf{s}) + \nu_i \cdot OPT$. More precisely, players are provided this additional utility if and only if they stick to their prescribed strategy under \mathbf{s} . A solution is said to be successfully stabilized if after the additional payments, no player wishes to deviate from strategy s_i . Our goal is to bound ν , the total payment that is required as a fraction of the optimum welfare, i.e., $\nu \cdot OPT = \sum_i \nu_i \cdot OPT$. Now suppose that for a given instance of the coordination game, the α -Approximate Equilibrium returned by our Algorithm of Theorem 4.5 has a social welfare of $\rho_\alpha \cdot OPT$ for $\rho_\alpha \leq 1$. Note that ρ_α is actually a function of the given instance which we omit as its meaning here is clear. Our second main result is the following.

Theorem 5.1. *For any given instance, the optimum solution can be stabilized by providing direct payments to players such that the total payment is at most a fraction $\nu \leq \frac{\rho_\alpha}{(\alpha-1)}$ times the social welfare of OPT .*

The above result indicates that if we run our algorithm from Theorem 4.5 and get a solution with low social welfare (ρ_α is small), then we can stabilize the optimum solution with a total payment that is approximately equal to $\rho_\alpha \cdot OPT$. We previously observed that when both γ and m are large, the solution returned by our algorithm may not be too efficient. In such cases, we can do much better by implementing the optimal solution and providing small incentives to players.

Proof. To be concrete, let us assume that we have run the Algorithm of Theorem 4.5 and obtained a solution whose social welfare is a fraction ρ_α times OPT for that given instance. Let s^* be the social optimum and by definition, $u(s^*) = A(s^*) + P(s^*)$, i.e., the components of the total welfare derived from players’ intrinsic preferences and that obtained from player relationships respectively. Let $X \subseteq N$ denote the players to whom payments have to be provided in s^* in order to prevent them from deviating and suppose that $\forall i \in X, k_i$, denotes their best-response strategy. Now how much additional utility $\nu_i OPT$ should we provide player i ? The minimum such incentive required must equal the additional utility that she would gain by performing best-response from OPT ,

i.e., $v_i OPT = u_i(k_i, s_{-i}^*) - u_i(s^*)$. The second term in the right hand side is bounded by the utility that players derive from the anchors alone. Therefore, we have,

$$\sum_{i \in X} (u_i(k_i, s_{-i}^*) - w_i^{s_i^*}) \geq \sum_{i \in X} (u_i(k_i, s_{-i}^*) - u_i(s^*)).$$

It suffices therefore to say that the term in the LHS of the above inequality serves as an upper bound for the total payments required to stabilize the optimum. Let us overload notation and represent $\sum_{i \in X} u_i(k_i, s_{-i}^*) = A(X) + P(X)$, where as usual the first term represents the utility coming from the intrinsic preferences and the second term the coordination utility. Now, we have $\nu OPT \leq A(X) + P(X) - \sum_{i \in X} w_i^{s_i^*}$.

Consider a strategy vector s' , where all players choose the same strategy, say k_0 , for some $1 \leq k_0 \leq m$. We know, $u(OPT) \geq u(s')$. We therefore, have $A(s^*) + P(s^*) \geq A(s') + P(s') \geq P(s') = P_T \geq P(s^*) + P(X)$. The final inequality comes from the fact that $P(X)$ and $P(OPT)$ are disjoint, i.e., the coordination utility that players would derive by deviating from OPT is not a part of OPT itself. But together $P(X) + P(OPT) \leq P_T$, which represents the total coordination utility derived when all players are choosing the same strategy. Removing $P(s^*)$ from both sides, we get, $A(s^*) \geq P(X)$. Finally we have,

$$\nu OPT \leq A(X) + P(X) - \sum_{i \in X} w_i^{s_i^*} \quad (2)$$

$$\leq A(X) + A(s^*) - \sum_{i \in X} w_i^{s_i^*} \quad (3)$$

$$= \sum_{i \in X} w_i^{k_i} + \sum_{i \in N} w_i^{s_i^*} - \sum_{i \in X} w_i^{s_i^*} \quad (4)$$

$$= \sum_{i \in X} w_i^{k_i} + \sum_{i \notin X} w_i^{s_i^*} \quad (5)$$

$$\leq A_T. \quad (6)$$

So we now have A_T , the sum of the utilities that each selfish player derives from choosing her favorite strategy as an upper bound for νOPT . We know from the proof of Theorem 4.5, that our α -Approximate Equilibria have a utility of at least $A_T(\alpha - 1)$. So we conclude $\nu OPT \leq \frac{\rho\alpha}{\alpha-1}$. \square

Putting this in perspective, if our 1.618-Approximate Equilibrium has a Social Welfare that is one-tenth of OPT , then we can stabilize the optimum by providing total payments no greater than sixteen percent of OPT . The bottom line for a central enforcer is that a little incentivizing goes a long way in ensuring socially efficient outcomes in coordination games. For every given instance, we can either apply the algorithm of Theorem 4.5 to compute a good quality approximate equilibrium or enforce even the optimum solution with a small total budget, thereby always ensuring high welfare, stable outcomes. Unfortunately, computing the optimum is NP-hard and therefore, we consider stabilizing other

high quality solutions that are a good approximation to OPT . In particular, we claim that the 0.8535 approximation to OPT provided by the polynomial-time algorithm in [28] can be stabilized with the same total payments as required for the optimum.

Corollary 5.2. *There exists a solution computable in polynomial time whose social welfare is at least 0.8535 times OPT , which can be stabilized by providing total payments ν no greater than $\frac{\rho_\alpha}{\alpha-1}$ times OPT .*

(*Proof Sketch*) We first take the 0.8535 approximation to the optimum returned by the Algorithm in [28]. Then we consider a solution where all players are present in k^* , where k^* is the strategy that maximizes the total intrinsic utility derived by players as we defined previously. Now, let \mathbf{s} be the solution among these two that has higher social welfare. We claim that \mathbf{s} can be stabilized with a total payment of $\frac{\rho_\alpha}{\alpha-1}$ times OPT . We do not go over the proof as it is exactly similar to the proof of Theorem 5.1. ■

5.1 Influencing Strong Equilibrium by strengthening relationships

Previously, we showed that although strong equilibrium may not exist in general, a 2-Approximate strong equilibrium always exists and can be computed. In this section, we look at α -Approximate Strong Equilibrium for $1 \leq \alpha \leq 2$ and show that these always exist in the presence of a little additional influence. Our influencing technique is quite different from our previous strategy of providing players direct incentives. Specifically, player i receives full utility from all the other players choosing the same strategy such that $w(i, j) > 0$ and a small fraction ω of the utility from players with whom they originally had $w(i, j) = 0$. From the perspective of viewing our game as a network game, we are essentially creating links between players. We only consider SCGs that obey the correlated coordination condition to eliminate uncertainty over the reward sharing.

Model: We first consider an equivalent formulation of the SCGs that obey the CC condition. Namely, let us associate two parameters with each player a_i and b_i such that when players i and j coordinate, player i receives a utility of $a_i b_j w(i, j)$ and player j receives $a_j b_i w(i, j)$. It is not hard to see that this is at least as general as the CC condition³. We add an additional constraint over the weights, namely $w(i, j) = 0$ or $w(i, j) = 1$ or $w(i, j) = -\infty$. The minus infinity weights can be viewed as conflict edges. In other words, players i and j can never choose the same strategy. Such a model with conflict edges was first considered in [26] to study the formation of gossip networks. Given such an SCG, we now define the ω extension of the game for $\omega \in [\frac{1}{2}, 1]$ to be the social coordination game with the same parameters except $w(i, j) = \omega$ where it was previously zero.

The incentivizing technique that we use here is quite intuitive. We let players derive full benefit from their neighbors present in the same group, but also give

³given $\gamma = (\gamma_1, \dots, \gamma_N)$ and $w(i, j)$ for all (i, j) , we simply replace $w(i, j)$ with $(\gamma_i + \gamma_j)w(i, j)$, set $a_i = \gamma_i \forall i$ and $b_i = 1 \forall i$.

them to a fraction ω utility from edges where $w(i, j) = 0$ originally. Formally, given a SCG, let $N(i)$ be all the j such that $w(i, j) = 1$ and $F(i)$ be all the j such that $w(i, j) = 0$. After incentivizing, in any given strategy \mathbf{s} , a player's utility is

$$u_i(\mathbf{s}) = \sum_{\substack{s_j = s_i \\ j \in N(i)}} a_i b_j + \omega \sum_{\substack{s_j = s_i \\ j \in F(i)}} a_i b_j$$

Our following result can be viewed as an extension of a result in [26], showing the existence of Approximate Strong Equilibria.

Proposition 5.3. *The ω -extension of any given SCG obeying the correlated coordination condition with weights as defined above admits a $\frac{1}{\omega}$ -Approximate Strong Equilibrium.*

We use the technique of lexicographic ordering that has been used for showing the existence of strong equilibrium in various guises in literature [23]. The appropriate definition for lexicographic ordering is present in the appendix.

Proof. Consider a m -dimensional vector $\pi : A \mapsto \mathbb{R}_+^m$ defined on each state \mathbf{s} such that the k^{th} component of the vector $\pi_k(\mathbf{s})$ is equal to $\sum_{i: s_i = k} b_i$. Recall that the set of allowed states A only includes strategy vectors where no two conflicted users are present in the same group.

Now consider the solution \mathbf{s}^* such that all the other solutions are sorted lexicographically smaller than \mathbf{s}^* , i.e. $\pi(\mathbf{s}) \preceq \pi(\mathbf{s}^*)$ for all \mathbf{s} . We claim that \mathbf{s}^* is a $\frac{1}{\omega}$ -approximate strong equilibrium. Consider any deviation by a group of players from \mathbf{s}^* to a new solution \mathbf{s} and let i be the player whose original strategy had the largest weight (smallest index) under the lexicographic ordering. Suppose i deviates from strategy k to k' , we claim that $\pi_k(\mathbf{s}^*) \geq \pi_{k'}(\mathbf{s})$. If this were not true, then the strategy vector \mathbf{s} would be lexicographically larger than \mathbf{s}^* . Now the utility that i received originally is

$$u_i(\mathbf{s}^*) = a_i \left(\sum_{\substack{s_j^* = s_i^* \\ j \in N(i)}} b_j + \omega \sum_{\substack{s_j^* = s_i^* \\ j \in F(i)}} b_j \right) \geq \omega a_i \sum_{\substack{j \neq i \\ s_j^* = k}} b_j = a_i \omega (\pi_k(\mathbf{s}^*) - b_i).$$

i 's utility after deviating is $u_i(\mathbf{s}) \leq a_i \sum_{s_j = k'} b_j = a_i (\pi_{k'}(\mathbf{s}) - b_i) \leq a_i (\pi_k(\mathbf{s}^*) - b_i) \leq \frac{1}{\omega} u_i(\mathbf{s}^*)$. This completes the proof of the theorem. \square

Corollary 5.4. *When $\omega = 1$, every instance of the ω -extension of SCGs obeying with the correlated coordination condition admits a Strong Equilibrium.*

6 Social Coordination with Complementarities

In this section, we generalize our model of social coordination to include scenarios where the benefits due to coordination may exhibit complementarities or supermodularities. So far, we have looked at incentivizing players and computing approximate equilibrium for a very specific model of social coordination

where the benefits are additively separable. However, it is not hard to see that most of our results extend to the case when utilities and welfare may be sub-modular in the set of players. We now generalize this further and consider SCGs wherein the benefits gained by coordinating with a set of players can be much larger than the utility gained by individually coordinating with each of them.

Model for Generalized Social Coordination Game. We consider a game with players belonging to a set N , a strategy set M and functions $w : 2^{N \cup M} \mapsto \mathbb{R}^+$ and $u_i : 2^{N \cup M} \mapsto \mathbb{R}^+$ for all players $i \in N$. Players still choose a single strategy and in any given state of the game \mathbf{s} , welfare due to all players X_k choosing strategy k is given by $w(k + X_k)$. This welfare is then distributed among the players, so we have

$$\sum_{i \in X_k} u_i(\mathbf{s}) = \sum_{i \in X_k} u_i(k + X_k) = w(k + X_k).$$

The total social welfare of the solution is given by $u(\mathbf{s}) = \sum_i u_i(\mathbf{s}) = \sum_{k \in M} w(k + X_k)$. Observe that we have made no assumptions on the functions w and u_i . Therefore, we may have that $w(S \cup T) \geq w(S) + w(T)$, which is not possible in our previous model. We only require that the total welfare generated is divided among the players.

There has always been considerable interest in utility-maximization games where either the resources or players exhibit complementarities. Situations of interest include combinatorial auctions wherein valuation functions can be supermodular in the set of items received, or group formation where people may possess specific sets of skills. However, there has not been much progress in this realm of algorithmic game theory as even the fundamental algorithmic problem of maximizing welfare in the presence of complements is considered *notoriously difficult*. Recent papers [1, 17, 29] have sought to overcome this obstacle by limiting the amount of supermodularity (or complementarity) present in the valuations in combinatorial auctions. With regard to full-information games such as ours, most similar to our model is the welfare-sharing game proposed by Bachrach et al. [6] who showed bounds on the Strong Price of Anarchy. Instead, here we focus on questions related to existence and computation of stable solutions.

We begin by showing that there exist instances of the Generalized SCG (with complementarities) that exhibit no stable solutions when the criterion considered is (Approximate) Nash Stability. Previously, we proved that all instances of SCGs without complementarity admit a 1.618-Approximate Equilibrium. However, when the player relationships become complementary, there exist instances which do not admit an α -Approximate Equilibrium for any $\alpha \geq 1$.

Proposition 6.1. *For any $c \geq 1$, there exists an instance of the Generalized Social Coordination Game which does not admit an α -Approximate Equilibrium for all $1 \leq \alpha \leq c$.*

Proof. The example that we use here is an extension of the triangle example from Proposition 3.1. Once again consider three players i, j, k and three strategies

1, 2, 3. We define the utility function for player i for different scenarios. The functions for players j and k are defined similarly. $u_i(1) = u_i(1+k) = c^2$, $u_i(1+j) = u_i(1+j+k) = c^3$, $u_i(2) = c$, $u_i(2+j) = u_i(2+j+k) = c^3$, $u_i(3) = u_i(3+k) = 0$, $u_i(3+j+k) = c$. Once again, it is not hard to verify that in any state at least one player can deviate and improve her utility by a factor c . Therefore, this instance admits no α -Approximate Equilibrium for any $\alpha \leq c$. \square

The non-existence of Approximate Nash Equilibrium for Generalized SCGs forces us to consider a more restricted model. We consider two forms of limited complementarities that have also been considered in other settings in literature. First, we look at hypergraph SCGs, a direct extension of the Network Game model mentioned in Section 2. We show some sufficient conditions for the existence of a Nash Equilibrium. We then move on to a model where the degree of supermodularity is bounded by some constant r . We adapt our *One-Shot* α -BR algorithm to obtain a $\Theta(r)$ -Approximate equilibrium.

6.1 Hypergraph Social Coordination Games

Once again, we view our SCG as a network game where the nodes include the players and the anchored nodes (strategies). There exist a set of hyperedges E , and each hyperedge has a weight w_e . The total welfare of a set S is given by $w(S) = \sum_{e \in S} w_e$, where the summation is over the hyperedges whose member vertices all belong to S . The welfare due a hyperedge is distributed among the vertices that comprise the hyperedge. For each player $i \in e$, we have a constant γ_i^e such that i 's utility due to the hyperedge is given by $\gamma_i^e w_e$. As is natural, the constants γ_i^e satisfy $\sum_{i \in e} \gamma_i^e = 1$. Such a model with a hypergraph induced on a set of items (instead of the players) was considered in [1] and they provided a r -Approximation Algorithm for welfare maximization where r is the maximum number of vertices present in a single hyperedge. We now extend our previous Correlated Coordination Condition for this Generalized SCG and show that Nash Equilibrium always exists for instances that obey the CC condition.

(Correlated Coordination Condition) A given instance of the SCG on hypergraphs is said to satisfy this condition if \exists a vector $\gamma = (\gamma_1, \dots, \gamma_N)$ such that $\forall e \in E$ with $w_e > 0$, we have $\gamma_i^e = \frac{\gamma_i}{\sum_{j \in e} \gamma_j}$ for all the vertices that belong to the hyperedge.

Proposition 6.2. *Hypergraph Social Coordination Games which obey the Correlated Coordination (CC) Condition admit a potential function. Therefore, best-response dynamics always converge to a Nash Equilibrium in such games.*

Proof. Given a social coordination game with weights $w_e, \forall e \in E$ and $\gamma = (\gamma_1, \dots, \gamma_N)$, that obeys the CC condition, we claim that the following is an inexact potential function for this game.

$$\Phi(s) = \sum_{\substack{e: \\ s_i = s_j \forall i \in e, j \in e}} \frac{w_e}{\sum_{j \in e} \gamma_j}.$$

Note that $\gamma_i = 0$ for all the anchored nodes as they do not receive any utility. The proof follows along the lines of the proof of Theorem 3.2 \square

6.2 Social Coordination Games with bounded supermodularity

We consider general SCGs where the welfare and utility may exhibit complementarities but we bound the amount of complementarity in each player's utility by a multiplicative factor $r \geq 1$. Specifically, we define a r -Supermodular SCG as a game where for all players i , we have $u_i(S \cup T) \leq r(u_i(S) + u_i(T))$ for any S and T . Informally this implies that even though the welfare due to a group of people coordinating is more than the sum of welfare due to smaller subgroups coordinating, the additional welfare due to complementarity is bounded by some parameter. The Social Coordination Games considered in Sections 1-5 satisfy this criterion with $r = 1$ and therefore, r -Supermodular SCGs are natural generalizations of linear or submodular Social Coordination Games.

Similar notions of restricted complementarities were studied in [29] and [17]. Both papers considered the optimization problem of welfare maximization and obtained approximation factors linear in the degree of supermodularity. We now consider the problem of computing an approximate equilibrium for SCGs with bounded supermodularity and show that every r -Supermodular SCG admits an approximate equilibrium whose stability factor is linear in the parameter r . It is an interesting open question whether this bound is tight or whether we can obtain stability factors sublinear in r .

Once again we turn to our *One-Shot α -BR* algorithm and adjust its parameter to obtain an $r + \epsilon$ -Approximate Equilibrium for some $\epsilon < 1$.

Theorem 6.3. *For any given instance of a r -Supermodular SCG, \exists an $\epsilon < 1$, such that running the One-Shot α -BR algorithm for $\alpha = r + \epsilon$ returns a $r + \epsilon$ -Approximate Nash Equilibrium.*

In other words, we can always compute a $\Theta(r)$ -Approximate Nash Equilibrium for any r -Supermodular SCG.

Proof. Let k_0 be the starting strategy for all players. Once again we assume that players deviate from k_0 in rounds, let $X_k(t)$ refer to the set of players whose strategy is k after t rounds. Suppose we allow players to deviate only when it improves their utility by a factor α or more (we will see what this factor is shortly). Let \mathbf{s} be the final solution. For every player i in \mathbf{s} whose strategy is k_0 , she cannot improve her utility by a factor more than α if she deviates (by definition). For any player who deviated to a strategy $k \neq k_0$ at some time t , if her best-response is k' , we then have

$$u_i(k', s_{-i}) = u_i(X'_k) \leq u_i(X'_k(t) \cup X_{k_0}(t)) \leq r(u_i(X'_k(t)) + u_i(X_{k_0}(t))).$$

But since k was i 's best-response at time t , we have $u_i(\mathbf{s}) \geq u_i(X'_k(t))$ and

$u_i(\mathbf{s}) \geq \alpha u_i(X_{k_0}(t))$. So we can finally conclude that

$$u_i(k', s_{-i}) \leq r(u_i(X'_k(t) + u_i(X_{k_0}(t))) \leq ru_i(\mathbf{s})(1 + \frac{1}{\alpha}).$$

The stability factors for the two sets of players are α and $r(1 + \frac{1}{\alpha})$ respectively. In the best possible case, they are both equal, giving us $\alpha = \frac{1}{2}(r + \sqrt{r(r+4)}) < r + 1$ for $r \geq 1$. \square

Acknowledgements

This work was supported in part by NSF awards CCF-0914782, CCF-1101495, CNS-1017932, and CNS-1218374.

References

- [1] Ittai Abraham, Moshe Babaioff, Shaddin Dughmi, and Tim Roughgarden. Combinatorial auctions with restricted complements. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 3–16. ACM, 2012.
- [2] Elliot Anshelevich, John Postl, and Tom Wexler. Assignment games with conflicts: Price of total anarchy and convergence results via semi-smoothness. *arXiv preprint arXiv:1304.5149*, 2013.
- [3] John Augustine, Ning Chen, Edith Elkind, Angelo Fanelli, Nick Gravin, and Dmitry Shiryayev. Dynamics of profit-sharing games. In *Proceedings of the Twenty-Second international joint conference on Artificial Intelligence-Volume Volume One*, pages 37–42. AAAI Press, 2011.
- [4] Haris Aziz, Felix Brandt, and Hans Georg Seedig. Computing desirable partitions in additively separable hedonic games. *Artificial Intelligence*, 195:316–334, 2013.
- [5] Yoram Bachrach, Edith Elkind, Reshef Meir, Dmitrii Pasechnik, Michael Zuckerman, Jörg Rothe, and Jeffrey S Rosenschein. The cost of stability in coalitional games. In *Proceedings of the 2nd International Symposium on Algorithmic Game Theory (SAGT)*, pages 122–134. Springer, 2009.
- [6] Yoram Bachrach, Vasilis Syrgkanis, Eva Tardos, and Milan Vojnovic. Strong price of anarchy and coalitional dynamics. *arXiv preprint arXiv:1307.2537*, 2013.
- [7] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Improved equilibria via public service advertising. In *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 728–737, 2009.

- [8] Suryapratim Banerjee, Hideo Konishi, and Tayfun Sönmez. Core in a simple coalition formation game. *Social Choice and Welfare*, 18(1):135–153, 2001.
- [9] Anand Bhalgat, Tanmoy Chakraborty, and Sanjeev Khanna. Approximating pure nash equilibrium in cut, party affiliation, and satisfiability games. In *Proceedings of the 11th ACM conference on Electronic Commerce (EC)*, pages 73–82. ACM, 2010.
- [10] Anna Bogomolnaia and Matthew O Jackson. The stability of hedonic coalition structures. *Games and Economic Behavior*, 38(2):201–230, 2002.
- [11] Simina Brânzei and Kate Larson. Coalitional affinity games and the stability gap. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 79–84, 2009.
- [12] Flavio Chierichetti, Jon Kleinberg, and Sigal Oren. On discrete preferences and coordination. In *Proceedings of the fourteenth ACM conference on Electronic Commerce (EC)*, pages 233–250. ACM, 2013.
- [13] George Christodoulou, Vahab S Mirrokni, and Anastasios Sidiropoulos. Convergence and approximation in potential games. In *Proceedings of the 18th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 349–360. Springer, 2006.
- [14] Michael Suk-Young Chwe. Communication and coordination in social networks. *The Review of Economic Studies*, 67(1):1–16, 2000.
- [15] Russell Cooper. *Coordination games*. Cambridge University Press, 1999.
- [16] Andreas Darmann, Edith Elkind, Sascha Kurz, Jérôme Lang, Joachim Schauer, and Gerhard Woeginger. Group activity selection problem. In *Proceedings of the 8th Intl. Workshop Internet & Network Economics (WINE)*, pages 156–169. Springer, 2012.
- [17] Uriel Feige and Rani Izsak. Welfare maximization and the supermodular degree. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 247–256. ACM, 2013.
- [18] Michal Feldman and Tami Tamir. Approximate strong equilibrium in job scheduling games. *Journal of Artificial Intelligence Research*, 36(1):387–414, 2009.
- [19] Moran Feldman, Liane Lewin-Eytan, and Joseph Seffi Naor. Hedonic clustering games. In *Proceedings of the 24th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 267–276. ACM, 2012.
- [20] Martin Gairing and Rahul Savani. Computing stable outcomes in hedonic games. In *Proceedings of the 3rd International Symposium on Algorithmic Game Theory (SAGT)*, pages 174–185. Springer, 2010.

- [21] Andrea Galeotti, Sanjeev Goyal, Matthew O Jackson, Fernando Vega-Redondo, and Leeat Yariv. Network games. *The Review of Economic Studies*, 77(1):218–244, 2010.
- [22] Sanjeev Goyal and Fernando Vega-Redondo. Network formation and social coordination. *Games and Economic Behavior*, 50(2):178–207, 2005.
- [23] Tobias Harks, Max Klimm, and Rolf H Möhring. Strong nash equilibria in games with the lexicographical improvement property. In *Proceedings of the 5th Intl. Workshop Internet & Network Economics (WINE)*, pages 463–470. Springer, 2009.
- [24] Martin Hoefer. *Cost sharing and clustering under distributed competition*. PhD thesis, University of Konstanz, 2007.
- [25] Matthew O Jackson and Alison Watts. On the formation of interaction networks in social coordination games. *Games and Economic Behavior*, 41(2):265–291, 2002.
- [26] Jon Kleinberg and Katrina Ligett. Information-sharing in social networks. *Games and Economic Behavior*, 82:702–716, 2013.
- [27] Michael Kosfeld. Economic networks in the laboratory: A survey. *Review of Network Economics*, 3(1), 2004.
- [28] Michael Langberg, Yuval Rabani, and Chaitanya Swamy. Approximation algorithms for graph homomorphism problems. In *Proceedings of the 9th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APRPOX)*, pages 176–187. Springer, 2006.
- [29] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 18–28. ACM, 2001.
- [30] Brendan Lucier and Renato Paes Leme. Gsp auctions with correlated types. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 71–80. ACM, 2011.
- [31] Dov Monderer and Lloyd S Shapley. Potential games. *Games and economic behavior*, 14(1):124–143, 1996.
- [32] Stephen Morris. Contagion. *The Review of Economic Studies*, 67(1):57–78, 2000.
- [33] Tim Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 513–522. ACM, 2009.

Appendix

A Proofs from Section 3

Proposition 3.3. The Price of Anarchy (PoA) for the Social Coordination Game over all instances is at most m .

Proof. **Definition** A payoff-maximization game is (λ, μ) -semi-smooth if there exists a mixed strategy σ_i for each player i such that for every outcome s

$$\mathbf{E}_\sigma \left[\sum_{i \in N} u_i(\sigma_i, s_{-i}) \right] \geq \lambda \cdot u(\mathbf{s}^*) - \mu \cdot u(\mathbf{s}).$$

It is easy to show that a game being semi-smooth implies all the same results from [33] as smoothness, including a bound of $\frac{1+\mu}{\lambda}$ for the price of total anarchy. We now show that our Social Coordination Game is $(\frac{1}{m}, 0)$ -semi-smooth for all instances. This immediately gives us an upper bound of m for the Price of Anarchy not just for Nash Equilibrium but for other more general solution concepts as well.

For any given strategy vector \mathbf{s} , for each player i , let σ_i be the mixed strategy where the player chooses each available strategy with equal probability $\frac{1}{m}$. Then we have,

$$\mathbf{E} \left[\sum_{i \in N} u_i(\sigma_i, s_{-i}) \right] = \sum_{1 \leq k \leq m} \frac{1}{m} (w_i^k + \sum_{s_j=k} \gamma_{ij} w(i, j)) = \frac{1}{m} (\sum_{1 \leq k} w_i^k + \sum_{j \in N} \gamma_{ij} w(i, j)).$$

The final inside the parenthesis on the right hand side represents the maximum utility that player i could receive in any given solution and thereby acts as an upper bound for her utility in the optimum solution. Summing the above equation for all players, we have, $\sum_{i \in N} \mathbf{E} [u_i(\sigma_i, s_{-i})] \geq \frac{1}{m} \cdot u(\mathbf{s}^*)$. We therefore conclude that our Social Coordination Game is $(\frac{1}{m}, 0)$ -semi-smooth. \square

Proposition 3.5. The price of stability is at most $(2 - \frac{1}{m})$ in social coordination games with symmetric relationships, i.e. $\gamma = 1$. Further, this bound is tight.

Proof. **Definition** A strategy vector \mathbf{s}^0 is said to satisfy the **Minimum Intrinsic Preference condition**(MIP) if the sum of utilities derived by players purely from their intrinsic preference for their respective strategy ($w_i^{s_i^0}$) is at least $\frac{1}{m}$ times the utility players derive from the strategies in any other solution.

For a strategy vector \mathbf{s}^0 that satisfies the MIP condition, we are guaranteed that $A(\mathbf{s}^0) \geq \frac{1}{m} A(\mathbf{s})$ for any given strategy vector \mathbf{s} .

Proposition .4. Let \mathbf{s}^0 be a solution that satisfies the Minimum Intrinsic Preference condition. Then, any best-response dynamics starting from \mathbf{s}^0 results in a Nash equilibrium \mathbf{s} such that the social welfare of \mathbf{s}^0 is at most $(2 - \frac{1}{m})$ times that of \mathbf{s} .

What this implies is that best-response dynamics from certain “good” starting points is guaranteed to terminate in Nash equilibria without much loss in social welfare. If the social optimum satisfies the MIP condition, then Proposition .4 tells us that there exists a Nash equilibrium whose social welfare is at least $\frac{1}{(2-\frac{1}{m})}$ times that of the optimum. We first prove the proposition and then show that the social optimum does indeed satisfy the MIP condition, which immediately proves Proposition 3.5.

(Proof of Proposition .4)

Let \mathbf{s} be the Nash equilibrium in which the best-response dynamics from a solution \mathbf{s}^0 that obeys the MIP condition terminates. Since the value of the potential function cannot decrease after any best-response sequence, we have $\Phi(\mathbf{s}) > \Phi(\mathbf{s}^0)$. This implies,

$$u(\mathbf{s}) = 2\Phi(\mathbf{s}) - A(\mathbf{s}) > 2\Phi(\mathbf{s}^0) - A(\mathbf{s}) = u(\mathbf{s}^0) + A(\mathbf{s}^0) - A(\mathbf{s}).$$

Since \mathbf{s}^0 satisfies the MIP condition, we know that $A(\mathbf{s}^0) \geq \frac{1}{m}A(\mathbf{s})$. This gives us

$$u(\mathbf{s}) > u(\mathbf{s}^0) - (1 - \frac{1}{m})A(\mathbf{s}) \quad (7)$$

$$\implies (2 - \frac{1}{m})u(\mathbf{s}) > u(\mathbf{s}^0), \quad (8)$$

where the last step comes from the fact that $A(\mathbf{s}) \leq u(\mathbf{s})$. This completes the proof of the claim.

It is easy to see that *OPT* does satisfy the MIP condition. First, we construct a solution \mathbf{s}_2 where all the players choose the same strategy, the one maximizing total utility. Since, there exist only m strategies, we know that the utility players derive from that strategy is at least $\frac{1}{m}$ times that of the sum of utilities all players derive from all strategies, which is in turn greater than A_T , i.e., we have $u(\mathbf{s}_2) \geq \frac{1}{m}A_T + P_T$. Since $u(OPT) \geq u(\mathbf{s}_2)$, we have $A(OPT) + P(OPT) \geq \frac{1}{m}A_T + P_T$. Note that $P(OPT)$ can never be greater than P_T since P_T includes all the edges between players in a given graph. This means that $A(OPT)$ has to be at least $\frac{1}{m}A_T$ for the inequality to hold. Now for any given solution \mathbf{s}' , $A(\mathbf{s}') \leq A_T$, which finally gives us the desired result.

$$A(OPT) \geq \frac{1}{m}A_T \geq \frac{1}{m}A(\mathbf{s}'),$$

for all \mathbf{s}' .

We note here that this gap between the social optimum and the stable solution maximizing welfare occurs due to the fact that edges between players count only once in the potential function but twice towards the social welfare. The worst case does indeed occur when players choose to follow their most preferred strategy selfishly as opposed to being with their neighbors. We now give a family of examples that illustrate this inherent inefficiency that comes with stability and in the process show that the bound of $2 - \frac{1}{m}$ is tight in the limiting case.

Example Consider an instance with m strategies and m players where each player has a different preferred strategy from which she derives a utility of

$w_i^i = r + \epsilon$. There is one player p such that $w(i, p) = r$ for all other players i . The only Nash equilibrium has a perfectly fragmented society where each player chooses her preferred strategy, achieving a total welfare of $mr + m\epsilon$, whereas all players choose the same strategy in the social optimum which has a welfare of $(2m - 1)r + \epsilon$. In the limiting case, as $\epsilon \rightarrow 0$, the ratio between the welfares approaches $2 - \frac{1}{m}$. \square

B Proofs from Section 4

Proposition 4.1. The following algorithm returns a $\sqrt{2}$ -Approximate Equilibrium for all instances with three strategies and no algorithm can give a better approximation factor over all instances as this ratio of 1.414 is tight.

1. Run *Algorithm 1* for $m = 2$ (for say strategies 1 and 2) and all players so that no player wants to deviate from one of those two strategies to the other.
2. While there exists a player, who can deviate from either strategy 1 or strategy 2 and improve her utility by a factor $\sqrt{2}$, allow her to deviate to strategy 3. Once that player deviates, run the algorithm for $m = 2$ (for strategies 1 and 2) so that players choosing these two strategies are stable with respect to each other.
3. Repeat step 2 until no player can to deviate to 3 and increase her utility by a factor $\sqrt{2}$.

Proof. Let the final solution be \mathbf{s} . We consider players choosing strategy 3 in rounds. Suppose i is the player choosing 3 at some time (round) t , then let s_{-i}^t be the strategy vector at time t indicating the strategies of all players other than i . Let $X_k(\mathbf{s})$ be the set of players whose strategy is k under \mathbf{s} . We have by definition,

$$\begin{aligned} u_i(3, s_{-i}^t) &\geq \sqrt{2}u_i(2, s_{-i}^t) \\ u_i(3, s_{-i}^t) &\geq \sqrt{2}u_i(1, s_{-i}^t) \end{aligned}$$

In other words, player i receives more utility by a factor $\sqrt{2}$ from strategy 3 than she does from either 1 or 2 at that time. Now observe that in the final solution, $X_1(\mathbf{s}) \cup X_2(\mathbf{s})$ cannot be more than the set of players present in strategies 1 and 2 at any time t before termination. Therefore, for player i , both $u_i(1, \mathbf{s}_{-i})$ and $u_i(2, (\mathbf{s}_{-i}))$ are bounded from above by $u_i(1, s_{-i}^t) + u_i(2, s_{-i}^t)$. So we finally have the maximum utility i can get by switching from strategy 3 to either 1 or 2 in \mathbf{s} ,

$$u_i(1, s_{-i}^t) + u_i(2, s_{-i}^t) \leq \frac{2u_i(3, s_{-i}^t)}{\sqrt{2}} \leq \sqrt{2}u_i(\mathbf{s}).$$

The last inequality comes from the fact that since the set of players in strategy 3 only increases in size, i 's final utility is at least as much as the utility she received when she deviated to 3. By definition, the algorithm terminates when players in strategies 1 and 2 cannot improve their utility by more than a factor $\sqrt{2}$ by deviating to 3, so the stability factor is trivially applicable for such players. Finally, players in strategy 1 do not wish to deviate to 2 and vice-versa by definition. This completes our proof. \square

Lemma 4.3. The *1-shot* α -BR algorithm, beginning with any starting strategy k_0 , returns a solution whose social welfare is at least $\frac{A_T}{\alpha}$

Proof. We use the same notation as before. Note that any player who performs her best-response gets utility at least $best(i) = \max_k w_i^k$, since she could always deviate to (or stay in) her preferred strategy. Her final utility is also at least $best(i)$ as player utilities do not decrease when more people choose the same strategy. Any player whose final strategy is k_0 cannot deviate and improve her utility by a factor more than α , by definition. Therefore, her utility has to be at least $\frac{best(i)}{\alpha}$. Suppose the final strategy vector is \mathbf{s} , we then have $u(\mathbf{s}) = \sum_i u_i(\mathbf{s}) \geq \sum_i \frac{best(i)}{\alpha} \geq \frac{A_T}{\alpha}$. \square

Lemma 4.4. The *one-shot* α -BR algorithm with \mathbf{s}_0 as the starting state (where all players choose a strategy k_0), results in a social welfare of at least $A(\mathbf{s}_0) + \frac{\alpha}{\gamma+1}P(\mathbf{s}_0)$, where γ is the Maximum Relationship imbalance.

Proof. Consider a deviation from some intermediate state \mathbf{s}' by player i , which improves her utility by a factor α . Then, the original social welfare decreases by (not counting the new utility yet) at most $w_i^{s'_i} + \sum_{s'_j=s'_i} (\gamma_{ij} + \gamma_{ji})w(i, j)$. But since, by definition of γ , we have $\gamma_{ji} \leq \gamma \cdot \gamma_{ij}$, we can now bound the decrease in utility by $w_i^{s'_i} + \sum_{s'_j=s'_i} (\gamma_{ij} + \gamma \cdot \gamma_{ij})w(i, j) = w_i^{s'_i} + (\gamma+1) \sum_{s'_j=s'_i} \gamma_{ij}w(i, j)$. The increase in social welfare is at least the new utility of player i which is a factor α times her old utility. Therefore this increase is at least, $\alpha(w_i^{s'_i} + \sum_{s'_j=s'_i} \gamma_{ij}w(i, j))$.

Therefore, in any deviation, for every unit of utility lost from a player's intrinsic preference, α units of utility are gained and for $(\gamma+1)$ units of utility from player coordination lost, α units are gained. Also observe that once a player deviates to a strategy, she cannot change after that, so players' utilities never decrease. It is not hard to verify that the maximum social welfare is lost when all the welfare lost comes from the player coordinations before a player switched from k_0 . Since the maximum utility due to coordinations in \mathbf{s}_0 is $P(\mathbf{s}_0)$, the resulting utility gain from the deviations is at least $\frac{\alpha P(\mathbf{s}_0)}{\gamma+1}$. In the worse case, the utility due to intrinsic preference that players in \mathbf{s}_0 is retained in \mathbf{s} , giving us the desired social welfare of $A(\mathbf{s}_0) + \frac{\alpha}{\gamma+1}P(\mathbf{s}_0)$. \square

Theorem 4.5. The following algorithm returns an α -Approximate Equilibrium for $\alpha \in [\phi, 2]$ whose social welfare is at least a fraction $\frac{\alpha - 1}{1 + \frac{(\gamma+1)}{\alpha}(\alpha - (1 + \frac{1}{m}))}$ of the optimum when $(\gamma + 1) \leq m\alpha$, and $\max(\frac{\alpha}{\gamma+1}, \frac{\alpha-1}{(m-1)+(\alpha-1)})$ times OPT otherwise.

Algorithm: “For a given α , run *One-shot α -BR* and *One-shot $\frac{1}{\alpha-1}$ -BR* with k^* as the starting strategy. Let the returned solutions be \mathbf{s}_1 and \mathbf{s}_2 . Return the solution among these two with greater social welfare.”

Proof. We begin by establishing that both the procedures in the algorithm result in a α -Approximate Equilibrium. Applying Lemma 4.2 to the solution returned by the *1-shot $\frac{1}{\alpha-1}$ -BR* procedure, we get that its stability factor is $\max(\frac{1}{\alpha-1}, \frac{1}{\frac{1}{\alpha-1}} + 1)$, that is $\max(\frac{1}{\alpha-1}, \alpha)$. \mathbf{s}_1 is, therefore, an α -approximate Nash equilibrium as $\alpha \geq \frac{1}{\alpha-1}$ for $\alpha \in [\phi, 2]$. Similarly, applying the lemma to \mathbf{s}_2 , we get that its stability factor is $\max(\alpha, \frac{1}{\alpha} + 1)$, which is also α for the given range. Since both \mathbf{s}_1 and \mathbf{s}_2 are α -Approximate Equilibrium, one can make the same conclusion for \mathbf{s} , since it is the best of the two solutions.

First, we claim that the total social welfare of the solution where all players choose the strategy k^* is at least $\frac{A_T}{m} + P_T$. Recall that A_T is the utility when each player chooses her favorite strategy. Since the total intrinsic utility derived when players choose k^* is at least as much as the intrinsic utility from any other strategy and since there are only a total of m strategies, we conclude that $A(\mathbf{s}_0) = \frac{A_T}{m}$, where \mathbf{s}_0 is the strategy under which all players choose k^* . $P(\mathbf{s}_0) = P_T$ since all players are present in the same strategy. Now, applying Lemma 4.3 to our algorithm’s α -BR procedure (where the starting state is k^*), we get that the final social welfare of the solution, $u(\mathbf{s})$, is at least $\frac{A_T}{m} + \frac{\alpha}{\gamma+1}P_T$. Applying Lemma 4.3 to our *One-shot $\frac{1}{\alpha-1}$ -BR* algorithm, we get that our computed solution has a social welfare of at least $(\alpha - 1)A_T$. Since our hybrid algorithm returns the best of these two solutions, we have

$$u(\mathbf{s}) \geq \max((\alpha - 1)A_T, \frac{A_T}{m} + \frac{\alpha}{\gamma+1}P_T).$$

Let us define $A_T = y(A_T + P_T)$, for some $y \leq 1$. Then, we get,

$$u(\mathbf{s}) \geq \max\left((\alpha - 1)y, \frac{\alpha}{\gamma+1} + y\left(\frac{1}{m} - \frac{\alpha}{\gamma+1}\right)\right) OPT \quad (9)$$

First, let’s assume that $\frac{\alpha}{\gamma+1} \geq \frac{1}{m}$. Then, observe that the first term in the parenthesis is increasing in y and the second term is decreasing. So, the maximum of the two terms is minimized when they are equal, i.e., $(\alpha - 1)y = \frac{\alpha}{\gamma+1} + y(\frac{1}{m} - \frac{\alpha}{\gamma+1})$ or $y = \frac{1}{1 + \frac{(\gamma+1)}{\alpha}(\alpha - (1 + \frac{1}{m}))}$. So we have $u(\mathbf{s})$ is always greater than $y(\alpha - 1)OPT = \frac{OPT}{1 + \frac{(\gamma+1)}{\alpha}(\alpha - (1 + \frac{1}{m}))}$.

Now, for the case when $\frac{\alpha}{\gamma+1} < \frac{1}{m}$, our results show that $u(\mathbf{s}) \geq \max(\frac{\alpha}{\gamma+1}, \frac{\alpha-1}{(m-1)+(\alpha-1)})$. The first term simply follows from equation 9 from which we note that when $\frac{\alpha}{\gamma+1} < \frac{1}{m}$, $u(\mathbf{s}) \geq \frac{\alpha}{\gamma+1}$. In order to show the second term, we prove a general result that is applicable for all instances, namely that the social welfare is bounded from below by $\approx \frac{1}{m}$.

Lemma .5. *The one-shot $\frac{1}{\alpha-1}$ -BR algorithm beginning from any starting state results in a social welfare which is at least a fraction $\frac{\alpha-1}{(m-1)+(\alpha-1)}$ times OPT.*

Proof. Consider the outcome of *One-Shot* $\frac{1}{\alpha-1}$ -BR algorithm. Any player who deviated from their initial strategy received at least a fraction $\frac{1}{m}$ times their utility in OPT as they performed best-response and their utility cannot decrease after that round. Players still present in the initial state k_0 in \mathbf{s} receive at least $\frac{\alpha-1}{(m-1)+(\alpha-1)}$ times their utility in OPT. To show this, observe that for any such i , $u_i(\mathbf{s}) \geq (\alpha-1)u_i(s'_i, s_{-i})$ for $s'_i \neq s_i$ and $u_i(\mathbf{s}) = u_i(s_i, s_{-i})$. Summing up these inequalities and the last equality, we have $((m-1) + (\alpha-1))u_i(\mathbf{s}) \geq (\alpha-1)u_i(\text{OPT})$, which gives our desired fraction. Now since every player is receiving at least the fraction $\frac{\alpha-1}{(m-1)+(\alpha-1)}$ of her utility in OPT, we have that the social of our solution is also the same fraction of the optimum solution welfare. This gives a lower bound on the social welfare of the solution returned by our algorithm for all cases and completes the proof of the theorem and sublemma. \square

\square

Proposition 4.8. There are instances of the symmetric Social Coordination Game with $m = 3$ groups where Strong equilibrium does not exist. However, whenever Strong equilibrium exists, its social welfare is no less than half of OPT when player relationships are symmetric.

Proof. The second part of the claim also applies for the symmetric SCG and technique we use is quite common in the Strong Price of Anarchy literature. Let \mathbf{s} be any Strong Nash equilibrium(SNE) and \mathbf{s}^* denote the optimal solution. Consider a group deviation from \mathbf{s} by the set of players whose strategy is k in OPT ($X_k(\mathbf{s}^*)$). Suppose these players deviate to strategy k from \mathbf{s} forming a new strategy vector \mathbf{s}' . Since \mathbf{s} is a SNE, in any group deviation there must be at least one player whose utility cannot increase. Let this player be i . Then, i must have received at least as much utility in \mathbf{s} as she does after deviating to k along with the players in $X_k(\mathbf{s}^*)$. This means that $u_i(\mathbf{s}) \geq u_i(\mathbf{s}') \geq u_i(\mathbf{s}^*)$, where the last inequality comes from the fact that in \mathbf{s}' , $X_k(\mathbf{s}^*) \subseteq X_k(\mathbf{s}')$.

Now consider a deviation by all players in $X_k(\mathbf{s}^*)$ to strategy k other than player i . Once again, \exists a player j such that j receives more utility from strategy s_j than from the set of players in $X_k(\mathbf{s}^*) - \{i\}$. So we get the inequality $u_j(\mathbf{s}) \geq u_j(\mathbf{s}^*) - \frac{w(i,j)}{2}$, since $\gamma_{ij} = \gamma_{ji} = 1$. If we again look at a deviation to strategy k by all players other than i and j , there must be one player whose utility after this deviation is not more than the utility before deviating in \mathbf{s} , giving us a

similar inequality as above. We therefore, repeat this process considering one less player each time, and sum over all the inequalities obtained. It is not hard to see that in the left hand side, we get $\sum_{i \in X_k(\mathbf{s}^*)} u_i(\mathbf{s})$ and the right hand side $\sum_{i; s_i^* = k} w_i^k + 0.5 \sum_{s_i^* = s_j^* = k} w(i, j)$, which is at least $\sum_{i \in X_k(\mathbf{s}^*)} \frac{u_i(\mathbf{s}^*)}{2}$. This is because each edge among the players of $X_k(\mathbf{s}^*)$ gets summed up only once (when one of the players leave). We now repeat process for all strategies k and set of players $X_k(\mathbf{s})$, and sum over all the sets

Now summing these inequalities over all players and all strategies, we get

$$\sum_i u_i(\mathbf{s}) \geq \sum_i \frac{u_i(\mathbf{s}^*)}{2},$$

which gives us our desired bound on the Strong Price of Anarchy. \square

Proposition 4.9. For $\alpha \geq 1$, the *One-shot α -BR* algorithm returns a $(1 + \alpha)$ -Approximate Strong Equilibrium.

Proof. We already know by Lemma 4.2 that any deviation by a single player cannot increase her utility by a factor greater than α . It is not immediately clear why this factor holds for group deviations. As mentioned previously, let V_1 refer to the players who deviated from the initial state k^* and V_2 be the players whose strategies under \mathbf{s} are k^* . We first claim that in any deviation by a set of players S that includes at least one player from V_1 , the first player in $V_1 \cap S$ to have originally deviated from k^* cannot improve her utility by more than a factor $(1 + \frac{1}{\alpha})$.

Consider any group deviation by players $S \subseteq V$ which contains at least one player from V_1 . Let the new strategy vector be \mathbf{s}_2 . Let i be the player in $S \cap V_1$ who deviated from k^* before all the other players in S , say at time t . Since all the other players in S had k^* as their strategy at time t , they must have belonged to $X_{k^*}^t$. Suppose i 's strategy in \mathbf{s}_2 is k , then the utility that i gets after deviating is at most $u_i(k, s_{-i}^t) + u_i(k^*, s_{-i}^t)$. But this is not greater than $u_i(s_i, s_{-i}^t) + \frac{1}{\alpha} u_i(s_i, s_{-i}^t)$, where the first term appears because s_i is i 's best response and the second is because the new utility after deviating is at least α times the old utility which was $u_i(k^*, s_{-i}^t)$. Therefore, we have $u_i(\mathbf{s}_2) \leq (1 + \frac{1}{\alpha}) u_i(\mathbf{s})$.

Now suppose $S \subseteq V_2$, i.e. $S \cap V_1 = \emptyset$. We know

$$u_i(\mathbf{s}_2) \leq u_i(k, s_{-i}) + u_i(k^*, s_{-i}) \quad (10)$$

for some $i \in S$, whose strategy under \mathbf{s}_2 is k . This is true because the all the new players present in strategy k under \mathbf{s}_2 must have all been players whose strategy in \mathbf{s} was k^* . But $\forall i \in V_2$, we know that the utility these players originally received when their strategy was k^* could not have improved by more than a factor α after deviation, if they had deviated individually. Applying this observation to equation 10, we have $u_i(\mathbf{s}_2) \leq u_i(k, s_{-i}) + u_i(k^*, s_{-i}) \leq \alpha u_i(\mathbf{s}) + u_i(\mathbf{s})$. Therefore, in any group deviation from \mathbf{s} , there is at least one

player whose utility increases by a factor no more than $(1 + \alpha)$ for the solution returned by the *One-shot α -BR* algorithm.

Specifically for $\alpha = 1$, we get a 2-Approximate Strong Equilibrium for every instance. Using the Sub-Lemmas of Theorem 4.5, we immediately get that this solution gives us a guarantee of at least half the social welfare of Theorem 4.5. \square

C Proofs from Section 5

Lexicographic Ordering

We borrow the following definition from [23].

Definition Let $\pi(\mathbf{s}_1)$ and $\pi(\mathbf{s}_2)$ be the potential vector for two states $\mathbf{s}_1, \mathbf{s}_2$ and denote by $\tilde{\pi}(\mathbf{s}_1), \tilde{\pi}(\mathbf{s}_2)$ be the sorted potential vectors respectively, derived by sorting the entries of the actual vector in a non-increasing order, that is $\tilde{\pi}_1(\mathbf{s}_1) \geq \dots \geq \tilde{\pi}_m(\mathbf{s}_1)$ and similarly for $\tilde{\pi}(\mathbf{s}_2)$. Then $\pi(\mathbf{s}_1)$ is strictly sorted lexicographically smaller than $\pi(\mathbf{s}_2)$ ($\pi(\mathbf{s}_1) \prec \pi(\mathbf{s}_2)$) if \exists an index k such that $\tilde{\pi}_l(\mathbf{s}_1) = \tilde{\pi}_l(\mathbf{s}_2) \forall l < k$ and $\tilde{\pi}_k(\mathbf{s}_1) < \tilde{\pi}_k(\mathbf{s}_2)$. Similarly, we say $\pi(\mathbf{s}_1)$ is sorted lexicographically smaller than $\pi(\mathbf{s}_2)$ ($\pi(\mathbf{s}_1) \preceq \pi(\mathbf{s}_2)$) when either $\pi(\mathbf{s}_1) \prec \pi(\mathbf{s}_2)$ or $\pi(\mathbf{s}_1) = \pi(\mathbf{s}_2)$.